


REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE 26-Aug-97	3. REPORT TYPE AND DATES COVERED Annual Report: 1 August 1996 - 30 August 1997	
4. TITLE AND SUBTITLE Automatic Construction of Accurate Models of Physical Systems		5. FUNDING NUMBERS C/G N00014-96-1-0720 PR 96PR05164-00	
6. AUTHOR(S) Elizabeth Bradley			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Colorado Department of Computer Science ECOT 717, CB 430 Boulder, CO 80309-0430		8. PERFORMING ORGANIZATION REPORT NUMBER None	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research Program Officer Michael Shneier ONR 311 Ballston Centre Tower One 800 N. Quincy Street, Arlington, VA 22217-5660		10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES None			
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for Public Release		12b. DISTRIBUTION CODE  Approved for public release Distribution Unlimited	
13. ABSTRACT (Maximum 200 words) This report describes recent progress in the development of PRET, a computer program that automates the process of system identification. Given hypotheses, observations, and specifications, PRET constructs an ordinary differential equation model of a target system with no other inputs or intervention from its user. The core of the program is a set of traditional system identification (SID) methods. A layer of artificial intelligence (AI) techniques built around this core automates the high-level stages of the identification process that are normally performed by a human expert.			
14. SUBJECT TERMS System Identification, Artificial Intelligence			15. NUMBER OF PAGES 6
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT Unclassified

NSN 7540-01-280-5500

DTIC QUALITY INSPECTED

Standard Form 298 (Rev.23-89)
 Prescribed by ANSI Std. Z39-18
 Z98-102

Automatic Construction of Accurate Models of Physical Systems

1997 Progress Report

25 August 1997

ONR #N00014-96-1-0720

PI: Elizabeth Bradley

University of Colorado

(303) 492-5355

lizb@cs.colorado.edu

Other Team Members:

Matthew Easley	graduate student
Apollo Hogan	professional research assistant
Agnes O'Gallagher*	Mathematician
Janet Rogers*	Computer Scientist
Reinhard Stolle*	graduate student
Laura Vidal	admin assistant

NIST

NIST

* Most of Stolle's funding is through another grant. Rogers and O'Gallagher are paid by NIST; they are devoting a few hours a month to consultations with our group as problems arise when we call their code from our programs.

FY97 Numerical Productivity Measures:

Refereed papers submitted but not yet published:	1
Refereed papers published:	4
Unrefereed reports and articles:	1
Books or parts thereof submitted but not yet published:	none
Books or parts thereof published:	none
Patents filed but not yet granted:	none
Patents granted:	none
Invited presentations:	3
Contributed presentations:	11
Honors received:	none
Prizes or awards received:	none
Promotions obtained:	none
Graduate students supported:	7 total/2 on this project
Post-docs supported:	2 total/1 on this project
Minorities supported:	1

1 Project Summary

Traditional system identification addresses the task of inferring a mathematical model of a system from observations of that system. A controls engineer might perform this task, in its most basic form, by choosing a power series and matching its coefficients, via some sort of regression, to a numerical time series from a sensor. The goal of this project is the development of a computer program, called PRET, that automates the entire system identification process by building an artificial intelligence (AI) layer around a set of traditional system identification techniques. This AI layer executes many of the high-level parts of the identification process that are normally performed by a human expert, using qualitative, symbolic, and geometric reasoning to perform *structural identification*, intelligently guiding the search to the proper type and form of equation. Unknown coefficients are then found with special, AI-adapted *parameter identification* techniques.

PRET works with ordinary differential equation (ODE) models, linear or nonlinear, in one variable or many. Its implementation combines traditional numerical analysis methods, such as nonlinear regression, with AI techniques like logic programming, symbolic and geometric reasoning, and constraint propagation. The inputs consist of specific information about an individual system, in three forms:

- the user's hypotheses about the physics involved
- observations, interpreted and described by the user, symbolically or graphically, in varying formats and degrees of precision
- physical measurements made directly and automatically on the system

To construct an ODE model from this information, PRET combines powerful mathematical formalisms, such as the link between the divergence of an ODE and the friction of the system that it describes, with domain-specific notions — such as force balances in mechanical systems — to allow the types of custom-generated approximations that are lacking in existing AI modeling programs. Two different sets of rules, both of which may easily be changed or augmented by the user, play critical roles in the model-building task. Domain-specific rules are used in the generate phase to combine hypotheses into models — a nontrivial task in a system with more than one degree of freedom, or a system in which physical effects couple to one another — while general rules about ODE properties are used by a custom first-order logic system to rule out models that conflict with the observations.

2 FY97 Research Accomplishments

- PRET efficiently and effectively solves a class of global optimization problems involved in the process of parameter estimation. Specifically, qualitative reasoning (QR) techniques are used to derive starting values for a nonlinear least-squares solver call. These good initial approximations guide the NLS solver away from local extrema in a complicated nonlinear regression landscape. QR techniques are also used to compute cutoff frequencies for digital filters in order to identify and remove noise.

Our paper on this topic was selected — along with seven others, out of a total of 76 — for a presentation of the highlights of the 1996 and 1997 QR workshops that is to be given at IJCAI 1997.

- PRET constructs models in a new¹ domain, viscoelastics, which has been the focus of state-of-the-art AI modeling programs by other QR groups. It was very easy for us to add this new domain because of the flexibility and extensibility designed into PRET's framework. This ease — and the comparisons to other work that the new domain implementation allowed — proved positive and satisfying.
- PRET automatically synthesizes model fragments from scratch using power-series techniques. This technique is ubiquitous in engineering, but nonexistent in AI modeling tools. It is not only powerful, but was also very easy to implement: for instance, a one- or two-term expansion in the first few derivatives of a state variable yield every form of force term found in freshman physics textbooks.
- PRET autonomously processes sensor data to construct observations of a target system. This involved solving part of control theory's *observer problem*: inferring internal system dynamics from incomplete output sensor data. We accomplish this using delay-coordinate embedding and then perform feature extraction on the resulting reconstructed dynamics with some simple geometric reasoning/computer vision techniques.

Our paper on these results, which was presented at IDA97 in London, was judged to be among the best papers of that conference.

- PRET now uses backward (SLD) reasoning². Backward reasoning is goal-directed, allowing invalid models to be ruled out more quickly, and it allows for smooth integration of the truth-maintenance module.
- We implemented a declarative meta-level control language with which the user can dynamically control the order in which inference rules are applied. The reason for this is that effective rule order choice often depends on current knowledge, and static rule ordering does not allow for the type of flexible control flow that exploits that knowledge.
- Finally, and perhaps most interesting of all, we have successfully used PRET on a real-world example — a radio-controlled car being used in a robotics system at the University of British Columbia. PRET's model was actually *better* than the one constructed by the project analyst in several senses, and the discrepancies helped the analyst refine his mental model of the system. More information on this example is included in section 4; powerpoint graphics will be enclosed in a separate email message.

3 FY98 Research Goals

While the items in the previous section represent significant successes — and large portions of the “to do” list from our 1996 Progress Report — the ultimate goal of this line of research

¹Previously, PRET could only solve classical mechanics problems.

²instead of breadth-first forward reasoning

is a tool that can construct models of high-dimensional black-box systems, drawn from any domain that admits ODE models, using only information that is observable from the ports of those systems. In order for PRET to attain the next level of this ambitious plan, a variety of significant improvements — the subject of the work to be performed under the remaining years of this contract — are necessary. These fall in four general areas, as listed below:

1. Expand PRET's capability to handle more-complex applications in more domains (electronics, chemical kinetics, etc.), more efficiently, using a richer rule set.
2. Use truth maintenance techniques to efficiently maintain partial proof trees that document the reasoning that has been performed. These proof trees avoid duplication of effort and will also be used to guide the generation of successor candidate models.
3. Develop and incorporate analysis *and control* techniques to actively and realistically exploit sensors *and actuators* to allow a true input-output approach to modeling, another idea that is relatively common in system identification but rare in AI.
4. Implement a graphical user interface to make the program easier for domain experts to use.

As mentioned above, we have made significant inroads on items 1-3, and we have successfully finished one of last year's action items — the power series techniques for automatic hypothesis generation. Item 4 is a new addition; our current interface makes demos rather painful. Specific subgoals for each item above are:

- Item 1:
 - We will continue adding domains to PRET. Electronics is currently first on the list; implementing this domain is nontrivial, as it will require us to rework and probably augment the syntax with which coordinates and connections are described to the program. (Bradley, Easley, Hogan, Stolle)
 - We will add rules as we test PRET on more examples in different domains. (Bradley, Easley, Hogan, Stolle)
 - We will begin collaborating with a group at MIT, headed by David Trumper in EECS, that is performing system identification on a commercial eyeglass lens-cutting machine, and attempt to reproduce their results automatically. (Bradley, Easley, Hogan, Stolle)
- Item 2:
 - We will finish the truth-maintenance system that allows knowledge reuse and provides explanations for the failure of models. (Hogan, Stolle)
 - We will finish the constraint propagation algorithm that will process constraint expressions on parameter values and state variables — for the purposes of (1) setting up the parameter estimation call (2) informing users about possible parameter values or (3) if the constraints cannot be satisfied, discarding the model. (Hogan, Stolle)

- Stolle will defend his Ph.D. thesis.
- Item 3:
 - We will finish setting up the hardware and software for the data acquisition and control system and complete the interface that allows PRET to interact with this equipment. (Easley; this was put on hold in FY97 by hardware problems.)
 - We will build and test another example system — probably a spring-mass oscillator on an air track. (Easley)
 - We will continue working out the syntax, semantics, and representation of the information PRET receives about the sensors and sends to the actuators. (Bradley, Easley)
 - We will begin to automate PRET’s output interactions. This involves automating much of what control theorists call *controllability* and *reachability*. (Bradley, Easley)
 - Easley will defend his Ph.D. proposal.
- Item 4:
 - We build a GUI and test it out, hopefully using someone from the MIT lens-cutting machine project as a test case. (Hogan)

4 Technical Transitions

As of last year, members of our group had verified PRET’s performance on some of the textbook examples used in the graduate system identification course offered by the Aerospace Engineering Department at the University of Colorado. Since then, PRET has been developed to the point where demos to outside people have become routine, which has elicited a variety of useful suggestions and directions (e.g., the GUI).

We have also used PRET on some real-world problems, including a radio-controlled car being used in a robotics system at the University of British Columbia. The input consisted of data files and email messages containing hypotheses about the physics (e.g., “ $\dot{\theta} = \rho v$ ”: the heading changes linearly with velocity) and qualitative observations of the R/C car (e.g., “it pulls to the left”). The goal was to duplicate the model found by a human analyst, but PRET’s model was actually *better*, in several senses. It was technically correct, but very different from his, and the disparities surprised him in interesting and useful ways. Specifically, the car’s initial velocity was modeled as negative, but the analyst knew that the car had started from rest — a fact that he *had not mentioned in the input information*. This violation of his intuition not only drew this piece of implicit knowledge (“ $v = 0$ at $t = 0$ ”) out into the explicit syntax of the PRET call, but also suggested a piece of physics that he had not noticed until then — there was a delay in the system. This delay is clearly visible on the powerpoint plots that accompany this report; if one adds a delay hypothesis to the PRET call, the program successfully models that behavior as well. From a cognitive science point of view, this interchange was particularly interesting, as the discrepancies helped the analyst refine his mental model of the system.

Another important feature of PRET's behavior is also visible on the plots. In this application, the engineers cared deeply about the *position* of the car, but the heading θ and the velocity v did not matter as much. In a case like this, an engineer — who wants the simplest possible ODE that models the important behavior — does not add terms to the model to account for glitches on an unimportant state variable, such as the bump in the velocity data. PRET behaves in exactly the same fashion: it is an engineer's modeling tool, not a scientific discovery system; because of this, it is specifically designed to identify and ignore unimportant details.

5 FY97 Publications and Submissions

All are full-length papers unless otherwise noted.

On this research project:

- E. Bradley and R. Stolle, "Automatic Construction of Accurate Models of Physical Systems," *Annals of Mathematics and Artificial Intelligence* 17:1-28 (1996).
- E. Bradley, A. O'Gallagher, and J. Rogers, "Global Solutions for Nonlinear Systems using Qualitative Reasoning," the *Eleventh International Workshop on Qualitative Reasoning about Physical Systems*, Cortona, Italy, May 1997. Currently in review for the *Annals of Mathematics and Artificial Intelligence*.
- E. Bradley and M. Easley, "Sensor Data Processing for Automated System Identification," the *Second International Workshop on Intelligent Data Analysis*, August 1997.
- R. Stolle and E. Bradley, "Opportunistic modeling," the *Third IJCAI Workshop on Engineering Problems in Qualitative Reasoning*, August 1997.

On other research projects:

- E. Bradley and J. Stuart, "Using Chaos to Generate Choreographic Variations," the *Fourth Experimental Chaos Conference*, August 1997. Currently in review for *Chaos*; see <http://www.cs.colorado.edu/~lizb/chaotic-dance.html>.
- J. Dixon, E. Bradley, and Z. Popović, "Nonlinear Time-Domain Analysis of Injection-Locked Microwave MESFET Oscillators," *IEEE Trans. on Microwave Theory and Technique*, 45:1050-1057 (1997).
- J. Hertzberg, C. Carlton, M. Schwieterman, E. Bradley, E. Davis, and M. Linne, "Splitting of Forced Elliptic Jets and Flames," the *Fourth International Microgravity Combustion Workshop*, NASA-Lewis Research Center, Cleveland Ohio, May 1997. [published abstract]

6 Online Information

Available on my home page, <http://www.cs.colorado.edu/~lizb/Home.html>, are copies of all of the papers listed above, along with short project descriptions intended to get graduate and undergraduate students interested in my research projects.